

CLAIMS

1 1. (Original) A fault tolerant computer system for executing one or more jobs on one or more
2 nodes, comprising,
3 a hierarchy of monitors for monitoring operations in the computer system including,
4 one or more first monitors for monitoring first operations and, for any
5 particular one of said first operations that fails, for restarting another instance
6 of said particular one of said first operations,
7 one or more second monitors for monitoring said first monitors and, if any
8 particular one of said first monitors fails, for restarting another instance of
9 said particular one of said first monitors.

2.

1 (Original) The system of Claim 1 wherein,
2 said one or more of said second monitors are monitored by at least one of said first
3 monitors and, if any particular one of said second monitors fails, said at least one of
4 said first monitors restarts another instance of said particular one of said second
5 monitors.

3.

1 (Original) The system of Claim 2 wherein one or more of said second monitors operates to commit
2 suicide if more than one of said another instance of said particular one of said second monitors is
3 restarted.

4.

1 (Original) The system of Claim 1 wherein,
2 said nodes operate to execute processes in a service unit, a communication unit and a
3 resource management unit.

5.

1 (Original) The system of Claim 1 wherein each of said nodes includes a computer having an
2 operating system, wherein pluralities of nodes form clusters and wherein each cluster has a
3 corresponding instantiation of said hierarchy of monitors for monitoring operations in the computer
4 system.

6.

1 (Original) The system of Claim 5 wherein each instantiation of said hierarchy of monitors includes,
2 a first instantiation of said one or more first monitors for monitoring first instantiation
3 operations and, for any particular one of said first instantiation operations that fails,

4 for restarting another instance of said particular one of said first instantiation
5 operations,
6 a second instantiation of said one or more second monitors for monitoring said first monitors
7 of said first instantiation and, if any particular one of said first monitors of said first
8 instantiation fails, for restarting another instance of said particular one of said first
9 monitors of said first instantiation.

7.

1 (Original) The system of Claim 5 including first and second instantiations and wherein,
2 said one or more of said second monitors of said second instantiation are monitored
3 by at least one of said first monitors of said first instantiation and, if any particular
4 one of said one or more of said second monitors of said second instantiation fails, for
5 restarting another instance of said particular one of said one or more of said second
6 monitors of said second instantiation.

8.

1 (Original) The system of Claim 1 wherein,
2 said second monitors maintain a record of particular ones of the first monitors that
3 are active and corresponding active particular ones of said first operations being
4 monitored by said particular ones of the first monitors.

9.

1 (Original) The system of Claim 8 wherein,
2 said second monitors use said record to ensure that active particular ones of said first
3 operations monitored by a failing one of said particular ones of the first monitors that
4 are active is monitored by a new instance of said failing one of said particular ones of
5 the first monitors that are active.

10.

1 (Original) The system of Claim 1 wherein said hierarchy of monitors includes,
2 one or more additional monitors for monitoring said first monitors or said second monitors,
3 and, if any particular one of said first monitors or said second monitors fails,
4 restarting another instance of said particular one of said first monitors or said second
5 monitors.

11.

1 (Original) The system of Claim 10 wherein said hierarchy of monitors includes,
2 one or more other monitors for monitoring said first monitors, said second monitors or said
3 additional monitors, and, if any particular one of said first monitors, said second
4 monitors or said additional monitors fails, restarting another instance of said

5 particular one of said first monitors, said second monitors or said additional
6 monitors.

12.

1 (Original) The system of Claim 1 wherein,
2 said first operations are jobs running on said nodes for providing services and, for
3 any particular one of said jobs that fails, one of said first monitors restarts another
4 instance of said particular one of said jobs.

13.

1 (Original) The system of Claim 12 wherein said jobs implement e-commerce transaction services.

14.

1 (Original) The system of Claim 12 wherein said jobs implement transaction services for financial
2 instruments.

15.

1 (Original) The system of Claim 12 wherein said first monitors are host agents for monitoring
2 operations of a plurality of jobs on a plurality of nodes where each job is monitored by only one of
3 said host agents.

16.

1 (Original) The system of Claim 12 wherein said first monitors are one or more agents operating on a
2 first level, each of said agents for monitoring operations of jobs on nodes where each job is
3 monitored by only one of said agents.

17.

1 (Original) The system of Claim 12 wherein,
2 said first monitors are one or more agents operating on a first level, each of said
3 agents for monitoring operations of jobs on nodes where each job is monitored by
4 only one of said agents, and
5 said one or more second monitors includes one or more local coordinators operating
6 on a second level where each local coordinator monitors one or more of said agents.

18.

1 (Original) The system of Claim 12 wherein,
2 said first monitors are one or more agents operating on a first level, each of said
3 agents for monitoring operations of jobs on nodes where each job is monitored by
4 only one of said agents, and wherein a particular one of said agents runs on a
5 particular one of said nodes where a job monitored by said particular one of said

6 agents runs.

19.

1 (Original) The system of Claim 12 wherein,

2 said first monitors are one or more agents operating on a first level, each of said
3 agents for monitoring operations of jobs on nodes where each job is monitored by
4 only one of said agents, and wherein a particular one of said agents runs on a
5 particular one of said nodes where a job monitored by said particular one of said
6 agents runs on other of said nodes than said particular one of said nodes.

20.

1 (Original) The system of Claim 12 wherein,

2 said first monitors are one or more agents operating on a first level, each of said
3 agents for monitoring operations of jobs on nodes where each job is monitored by
4 only one of said agents, and wherein a particular one of said agents runs on a
5 particular one of said nodes where a job monitored by said particular one of said
6 agents runs,
7 said second monitors are one or more local coordinators operating on a second level,
8 each of said local coordinators for monitoring operations of agents, and wherein a
9 particular one of said local coordinators runs on a particular one of said nodes where
10 an agent monitored by said particular one of said local coordinators runs.

21.

1 (Original) The system of Claim 12 wherein,

2 said first monitors are one or more agents operating on a first level, each of said
3 agents for monitoring operations of jobs on nodes where each job is monitored by
4 only one of said agents, and wherein a particular one of said agents runs on a
5 particular one of said nodes where a job monitored by said particular one of said
6 agents runs,
7 said second monitors are one or more local coordinators operating on a second level,
8 each of said local coordinators for monitoring operations of agents, and wherein a
9 particular one of said local coordinators runs on a particular one of said nodes other
10 than where an agent monitored by said particular one of said local coordinators runs.

22.

1 (Original) The system of Claim 12 wherein,

2 said first monitors are one or more agents operating on a first level, each of said
3 agents for monitoring operations of jobs on nodes where each job is monitored by
4 only one of said agents,
5 said second monitors are one or more local coordinators operating on a second level,

6 each of said local coordinators for monitoring operations of agents,
7 and wherein said hierarchy of monitors includes,
8 one or more third monitors for monitoring said one or more second monitors and, for
9 any particular one of said second monitors that fails, restarting another instance of
10 said particular one of said second monitors, and wherein a particular one of said third
11 monitors that monitors said particular one of said second monitors runs on a different
12 node than a node where said particular one of said second monitors runs.

23.

1 (Original) The system of Claim 22 wherein said hierarchy of monitors includes,
2 one or more fourth monitors for monitoring said one or more third monitors and, for
3 any particular one of said third monitors that fails, restarting another instance of said
4 particular one of said third monitors, and wherein a particular one of said fourth
5 monitors that monitors said particular one of said third monitors runs on a different
6 node than a node where said particular one of said third monitors runs.

24.

1 (Original) The system of Claim 12 wherein,
2 said first monitors are one or more agents operating on a first level, each of said
3 agents for monitoring operations of jobs on nodes where each job is monitored by
4 only one of said agents,
5 said second monitors are one or more local coordinators operating on a second level,
6 each of said local coordinators for monitoring operations of agents,
7 and wherein said hierarchy of monitors includes,
8 one or more third monitors for monitoring said one or more second monitors and, for
9 any particular one of said second monitors that fails, restarting another instance of
10 said particular one of said second monitors, and wherein a particular one of said third
11 monitors that monitors said particular one of said second monitors runs on a node
12 where said particular one of said second monitors runs.

25.

1 (Original) The system of Claim 24 wherein said hierarchy of monitors includes,
2 one or more fourth monitors for monitoring said one or more third monitors and, for
3 any particular one of said third monitors that fails, restarting another instance of said
4 particular one of said third monitors, and wherein a particular one of said fourth
5 monitors that monitors said particular one of said third monitors runs on a node
6 where said particular one of said third monitors runs.

26.

1 (Original) The system of Claim 1 wherein said hierarchy of monitors includes,
2 one or more third monitors for monitoring said one or more second monitors and, for
3 any particular one of said second monitors that fails, restarting another instance of
4 said particular one of said second monitors.

27.

1 (Original) The system of Claim 26 wherein one or more of said second monitors operates to commit
2 suicide if more than one of said instance of said particular one of said second monitors is restarted.

28.

1 (Original) The system of Claim 26 wherein said one or more third monitors run on different ones of
2 said nodes than ones of said nodes on which said second monitors run.

29.

1 (Original) The system of Claim 26 wherein said hierarchy of monitors includes,
2 one or more fourth monitors for monitoring said one or more third monitors and, for
3 any particular one of said third monitors that fails, restarting another instance of said
4 particular one of said third monitors.

30.

1 (Original) The system of Claim 29 wherein said one or more fourth monitors run on different ones
2 of said nodes than ones of said nodes on which said third monitors run.

31.

1 (Original) The system of Claim 29 wherein said one or more fourth monitors run on ones of said
2 nodes which are the same as ones of said nodes on which said third monitors run.

32.

1 (Original) The system of Claim 29 wherein one or more of said third monitors operates to commit
2 suicide if more than one of said instance of said particular one of said third monitors is restarted.

33.

1 (Original) The system of Claim 1 having a resource management unit including a load-balancing for
2 distributing jobs among said nodes.

34.

1 (Original) The system of Claim 1 having a resource management unit including a persistent storage
2 unit.

35.

1 (Original) The system of Claim 1 having a resource management unit including an interface unit.

36.

1 (Original) The system of Claim 1 wherein,
2 each of said nodes includes a plurality of computers each having an operating system.

37.

1 (Original) The system of Claim 1 having a plurality of clusters of said nodes, each cluster having a
2 corresponding instantiation of said hierarchy of monitors for monitoring operations in the computer
3 system.

38.

1 (Original) The system of Claim 37 wherein,
2 each of said clusters of nodes operates to execute processes organized into a service
3 unit, a communication unit and a resource management unit.

39. (Original) The system of Claim 37 wherein,

said clusters of nodes are organized into groups, each group having one or more of
said clusters.

40.

1 (Original) The system of Claim 37 wherein,
2 a first one of said groups is located at a geographic location remote from a second
3 one of said groups and said first one of said groups is connected to said second one of
4 said groups by one or more networks.

41.

1 (Original) The system of Claim 37 wherein,
2 a first one of said groups is organized to execute on one subset of data and a second
3 one of said groups is organized to execute on another subset of data.

42.

1 (Original) The system of Claim 37 wherein,
2 a first one of said groups is organized to execute on one subset of data and a second
3 one of said groups is organized to provide backup for said one subset of data.

43.

1 (Original) The system of Claim 1 wherein,
2 said first operations are jobs running on said nodes for providing services,
3 said first monitor senses one or more conditions that can cause any particular one of
4 said jobs to fail whether or not said particular one of said jobs has actually failed,
5 one of said first monitors terminates said particular one of said jobs and restarts
6 another instance of said particular one of said jobs.

44.

1 (Original) The system of Claim 43 wherein,
2 said one of said first monitors that terminates said particular one of said jobs restarts
3 said another instance of said particular one of said jobs in an environment where said
4 one or more conditions are not present.

45.

1 (Original) The system of Claim 43 wherein,
2 said one of said conditions is a node failure and said another instance of said
3 particular one of said jobs is started on a different non-failing node.

46.

1 (Original) The system of Claim 43 wherein,
2 said one of said conditions is a job failure and said another instance of said particular
3 one of said jobs is started as a new instance of said job.

47.

1 (Original) The system of Claim 46 wherein,
2 said another instance of said particular one of said jobs is started as a new instance of said job
3 on a node the same as a node on which said particular one of said jobs was running.

48.

1 (Original) The system of Claim 46 wherein,
2 said another instance of said particular one of said jobs is started as a new instance of said job
3 on a new node different from a node on which said particular one of said jobs was running.

49.

1 (Original) The system of Claim 1 wherein each of said nodes includes a computer and wherein new
2 ones of said nodes are added to the system without disturbing the operations of other of said nodes in
3 the computer system and wherein jobs are assigned dynamically to said new ones of said nodes.

50.

1 (Original) The system of Claim 1 wherein each of said nodes includes a computer and wherein ones
2 of said nodes are removed from the system without disturbing the operations of other of said nodes
3 in the computer system and wherein particular jobs are reassigned dynamically to other of said nodes
4 in the computer system.

51.

1 (Original) The system of Claim 1 wherein each of said nodes includes a computer of one type and
2 wherein new ones of said nodes are added to the system including upgraded computers of a different
3 type without disturbing the operations of other of said nodes in the computer system and wherein
4 jobs are assigned dynamically from said other of said nodes to said new ones of said nodes to
5 provide dynamic upgrade of said system without stopping said particular jobs.

52.

1 (Original) The system of Claim 1 wherein pluralities of nodes form clusters and wherein particular
2 ones of said clusters are assigned for processing particular jobs at particular times and wherein other
3 ones of said clusters are assigned for processing said particular jobs at other times.

53.

1 (Original) The system of Claim 52 wherein said particular times and said other times are follow-the-
2 sun times.

54.

1 (Original) The system of Claim 1 wherein a delay time is controlled before the restart of a job.

55.

1 (Original) The system of Claim 1 wherein a delay time is controlled before the restart of a job. An
2 interface that allows humans to monitor the health of the system and to log statistics about uptime of
3 each component in the system.

56.

1 (Original) The system of Claim 1 wherein a delay time is applied before said restarting another
2 instance of said particular one of said first operations.

57.

1 (Original) The system of Claim 1 wherein in said hierarchy of monitors,
2 said one or more of said second monitors are monitored by at least one of said first
3 monitors and, if any particular one of said second monitors fails, said at least one of
4 said first monitors, after a first delay time, restarts another instance of said particular
5 one of said second monitors on a node other than a node on which said particular one
6 of said second monitors failed.

58.

1 (Original) The system of Claim 57 wherein,
2 if more than one instance of said another instance of said particular one of said
3 second monitors is restarted, all but one instance of said another instance of said
4 particular one of said second monitors commits suicide.

59.

1 (Original) The system of Claim 57 wherein said hierarchy of monitors includes,
2 one or more additional monitors for monitoring said first monitors and said second monitors,
3 and, if any particular one of said first monitors or said second monitors fails,
4 restarting, after a second delay time, another instance of said particular one of said
5 first monitors or said second monitors.

60.

1 (Original) The system of Claim 59 wherein,
2 if more than one of instance of said another instance of said particular one of said
3 first monitors or said second monitors is restarted, all but one instance of said another
4 instance of said particular one of said first monitors or said second monitors operates
5 to commit suicide.

61.

1 (Original) The system of Claim 58 wherein said hierarchy of monitors includes,
2 one or more other monitors for monitoring said first monitors, said second monitors and said
3 additional monitors, and, if any particular one of said first monitors, said second
4 monitors or said additional monitors fails, restarting, after a third delay time, another
5 instance of said particular one of said first monitors, said second monitors or said
6 additional monitors.

62.

1 (Original) The system of Claim 61 wherein,
2 if more than one instance of said another instance of said particular one of said first
3 monitors, said second monitors or said additional monitors is restarted, all but one
4 instance of said another instance of said particular one of said first monitors, said
5 second monitors or said additional monitors operates to commit suicide.

63.

1 (Original) The system of Claim 1 wherein,
2 said first operations are jobs running on said nodes for providing services where a
3 particular first one of said jobs associated with a first customer is running on a

4 particular first node and a particular second one of said jobs associated with a second
5 customer is running on said particular first node.

64.

1 (Original) The system of Claim 1 wherein,

2 said first operations are jobs running on said nodes for providing services where a
3 particular first one of said jobs associated with a first customer is running on a
4 particular first node and a particular second one of said jobs associated with a second
5 customer is running on a particular second node whereby said first customer job is
6 isolated from said second customer job.

65.

1 (Original) The system of Claim 1 wherein,

2 said first operations are jobs running on said nodes for providing services where,
3 particular first ones of said jobs are associated with a first customer with one
4 of said particular first ones of said jobs running on a particular first node and
5 with another one of said particular first ones of said jobs running on a
6 particular other node;

7 particular second ones of said jobs are associated with a second
8 customer with one of said particular second ones of said jobs running on a
9 particular second node and with another one of said particular second ones
10 of said jobs running on said particular other node.

66.

1 (Original) The system of Claim 1 including transaction initiators for starting said first operations as
2 one or more jobs to initiate a transaction in a service.

67.

1 (Original) The system of Claim 1 including transaction processors for starting said first operations
2 as one or more jobs to process a transaction in a service.

68.

1 (Original) The system of Claim 1 including,

2 transaction initiators for starting first ones or more of said first operations as one or
3 more first jobs on a first node to initiate a transaction in a service;
4 transaction processors for starting other ones or more of said first operations as one
5 or more other jobs on another node to process said transaction in said service.

69.

1 (Original) The system of Claim 1 including,
2 transaction initiators for starting first ones or more of said first operations as one or
3 more first jobs on a first node to initiate a transaction in a service;
4 transaction processors for starting other ones or more of said first operations as one
5 or more other jobs on another node to process said transaction in said service.

70.

1 (Original) The system of Claim 1 including,
2 transaction initiators for starting first ones or more of said first operations as one or
3 more first jobs on a first node to initiate a transaction in a service;
4 transaction processors for starting other ones or more of said first operations as one
5 or more other jobs on said first node to process said transaction in said service.

71.

1 (Original) In a fault tolerant computer system operating to execute one or more jobs on one or more
2 nodes where the computer system includes a hierarchy of monitors for monitoring operations in the
3 computer system, the method comprising,
4 monitoring first operations with one or more first monitors and, for any particular one
5 of said first operations that fails, restarting another instance of said particular one of
6 said first operations,
7 monitoring said first monitors with one or more second monitors and, if any
8 particular one of said first monitors fails, restarting another instance of said particular
9 one of said first monitors.

72.

1 (Original) The method of Claim 71 wherein,
2 monitoring said one or more of said second monitors with at least one of said first
3 monitors and, if any particular one of said second monitors fails, restarting with said
4 at least one of said first monitors another instance of said particular one of said
5 second monitors.

73.

1 (Original) The method of Claim 2 wherein one or more of said second monitors operates to commit
2 suicide if more than one of said another instance of said particular one of said second monitors is
3 restarted.